

SEO · SEARCH VISIBILITY SPECIALIST

SEO, AEO and GEO Visibility Audit

Three retrieval systems, three different units of success, and one column that most of the industry is currently leaving out: how much of this you actually control.

THE DECISION THIS TEMPLATE PROTECTS

Where this brand is genuinely findable, across three surfaces that behave differently.

ROLE	COURSE PLACEMENT	REVISION
SEO · relevant to Content and Manager	Future course · Search Visibility Systems	Version 1.0 Current as of 19 August 2026

Audit each surface separately. They are not three names for the same activity, and treating them as one is why so much current advice is confused. Then tag every finding for **controllability** before you prioritise anything. A recommendation you cannot influence is a hypothesis, and it should be labelled as one when you present it.

BEFORE YOU START, YOU NEED

- Search Console and analytics access.
- A list of the twenty queries that matter commercially, not the two hundred that have volume.
- A way to check how AI assistants currently describe this brand and category. Manual checks are acceptable and currently more honest than most tooling.

01 The three surfaces, and why they are not the same thing

SURFACE	WHAT IS RETRIEVED	WHAT SUCCESS IS	HOW MUCH YOU CONTROL
SEO · classic search	A page, ranked among others	Position, click, session	Substantial. Technical, content, structure, links.

SURFACE	WHAT IS RETRIEVED	WHAT SUCCESS IS	HOW MUCH YOU CONTROL
AEO • answer engines and answer features	A passage, extracted to answer directly	Being the source of the answer, often with no click	Moderate. Passage structure, directness, schema, factual clarity.
GEO • generative and assistant retrieval	An entity and the claims attached to it, synthesised	Being named, recommended, and described correctly	Indirect. Consistency of claims across the web, third-party corroboration, crawlability.

The practical consequence: an AEO win can reduce traffic while increasing influence, and a GEO win may never appear in any analytics property you own. Have that conversation with the client **before** it happens, not in the month the graph goes down.

02 Surface one • Classic search baseline

QUERY	INTENT	CURRENT POSITION	PAGE	COMMERCIAL VALUE

WHICH OF THESE QUERIES STILL PRODUCE CLICKS, AND WHICH ARE ABSORBED BY ANSWER FEATURES

03 Surface two • Passage readiness

An answer engine extracts a passage. A passage that only makes sense in the context of the paragraphs around it cannot be extracted, however well the page ranks.

CHECK	PRESENT?	WHERE IT FAILS
Direct answers stated in the first two sentences under a heading		
Headings phrased as the questions people actually ask		
Definitions that stand alone without the surrounding paragraph		
Facts stated with a date and a source		
Tables and lists for genuinely comparative or enumerable content		
No critical claim buried mid-paragraph in a long block		

04 Surface three · Entity consistency

Generative systems assemble a description of the brand from many sources. Where those sources disagree, the description drifts. This is the audit almost nobody runs.

CLAIM ABOUT THE BRAND	HOW WE STATE IT	HOW THIRD PARTIES STATE IT	CONSISTENT?
What the company does			
Where it operates			
Who founded or leads it			
Category it belongs to			
Who it serves			
What it is known for			

HOW THE MAJOR ASSISTANTS CURRENTLY DESCRIBE THIS BRAND

Ask each one the same three questions and record the answers verbatim, with the date. Repeat quarterly. This is the only baseline you will have.

05 Machine readability

ITEM	STATUS	ACTION
Structured data present and valid		
Organisation and person entities marked up		
Author entities with credentials		
Dates: published and meaningfully updated		
Crawler access for AI user agents: decided deliberately		
llms.txt or equivalent, if we choose to publish one		
Content reachable without JavaScript execution		

Blocking AI crawlers is a legitimate strategic choice with a real cost. Make it deliberately and write down who made it. Discovering in month six that a developer blocked everything is a common and expensive surprise.

06 Third-party corroboration

Being described correctly by other people is the closest thing to a lever on generative retrieval that exists today.

SOURCE TYPE	DO WE APPEAR?	ACCURATE?	ACTION
Industry directories and listings			
Review platforms			

SOURCE TYPE	DO WE APPEAR?	ACCURATE?	ACTION
Trade and news coverage			
Community forums and discussion			
Comparison and roundup pages			
Reference sites and knowledge bases			

07 Prioritisation by controllability

This is the section that makes the audit honest. Sort every finding into one of three bands, then prioritise inside each band by business impact.

FINDING	SURFACE	CONTROLLABLE · INFLUENCEABLE · NOT CONTROLLABLE	BUSINESS IMPACT	PRIORITY

There is a market currently selling confident causal claims about systems nobody outside the model providers can observe. Refusing to make those claims costs a little credibility with buyers who want certainty, and earns a great deal with the ones who will still be your clients in three years.

08 Measurement plan, including what cannot be measured

WHAT WE WILL REPORT	HOW	CONFIDENCE	CADENCE
Ranked positions on priority queries			
Clicks and impressions			
Appearance in answer features			
Assistant descriptions of the brand			
Third-party mention accuracy			

WHAT WE ARE EXPLICITLY NOT ABLE TO MEASURE, AND WHAT WE WILL SAY ABOUT IT

JUDGMENT PROMPTS

1. If traffic halved but the brand became the source cited in AI answers, would this client call that a win, and have you had that conversation yet?
2. Which of your recommendations is a causal bet rather than an observed relationship, and is it labelled as one?
3. What would you stop doing if classic search traffic became a minority of discovery for this category?
4. Which finding is in the controllable band but has been ignored for months because it is unglamorous?

WHAT THIS TEMPLATE WILL NOT DO FOR YOU

- It will not tell you how any generative system ranks or selects sources. Nobody outside those companies can tell you that, and anyone who says otherwise is selling something.
- It has a short half-life. Check the review date at the foot of this page before relying on the surface definitions.
- It does not cover content quality, only its retrievability. A well-structured page that says nothing will be extracted and ignored.
- Manual assistant checks are a small, non-representative sample. Treat them as a baseline to watch, not as a measurement.

THE DEFENSIBLE OUTPUT

A three-surface audit with every finding tagged for controllability, a prioritised action list, and a measurement plan that states in writing what cannot currently be measured.

AFactor | An online-first institution for career transformation and future industry leadership.

This template is published open. Copy it, adapt it, use it on client work. If you pass it on, keep this line on it.

S1 · Version 1.0 · Current as of 19 August 2026 · Review cycle: every 90 days